arXiv:1607.05956v1 [cs.LO] 18 Jul 2016

Occam's Razor Applied to the Petri Net Coverability Problem

Thomas Geffroy, Jérôme Leroux, and Grégoire Sutre

Univ. Bordeaux & CNRS, LaBRI, UMR 5800, Talence, France

Abstract. The verification of safety properties for concurrent systems often reduces to the *coverability* problem for Petri nets. This problem was shown to be ExpSpace-complete forty years ago. Driven by the concurrency revolution, it has regained a lot of interest over the last decade. In this paper, we propose a generic and simple approach to solve this problem. Our method is inspired from the recent approach of Blondin, Finkel, Haase and Haddad [3]. Basically, we combine forward invariant generation techniques for Petri nets with backward reachability for well-structured transition systems. An experimental evaluation demonstrates the efficiency of our approach.

1 Introduction

Context. The analysis of concurrent systems with unboundedly many processes classically uses the so-called *counter abstraction* [12]. The main idea is to forget about the identity of each process, so as to make processes indistinguishable. Assuming that each process is modeled by a finite-state automaton, it is then enough to count, for each state q, how many processes are in state q. The resulting model is a Petri net, with no a priori bound on the number of tokens. The verification of a safety property on the original concurrent system (e.g., mutual exclusion) translates into a *coverability* question on the Petri net: Is it possible to reach a marking that is component-wise larger than a given marking?

Related work. Karp and Miller [14] proved in 1969 that coverability is decidable (but their algorithm is not primitive recursive), Lipton showed that it requires at least exponential space [15], and Rackoff showed that it only requires exponential space [17]. Despite these somewhat negative results, and driven by the concurrency revolution, the coverability problem has regained a lot of interest over the last decade. Recent efficient approaches include target set widening [13] and structural analysis mixed with SMT solving [7,3]. We believe that the time is ripe to experiment with new ideas and prototypes for coverability, and to apply them to real-world concurrent systems.

Our work builds notably on [3], which proposes a new approach to the coverability problem and its implementation. The approach of [3] is conceptually simple and exploits recent advances in the theory of Petri nets as well as the power of modern SMT-solvers. In a nutshell, they leverage recent results on coverability in continuous Petri nets [9] to over-approximate coverability under the

standard semantics: any configuration that is not coverable in a continuous Petri net is also not coverable under the standard semantics. This observation is then exploited inside a backward-coverability framework [1].

Our contribution. We present a generic backward coverability algorithm that relies on downward-closed (forward) invariants to prune the exploration of the state space. Our algorithm is in fact a family of algorithms parametrized by downward-closed invariants. It generalizes the algorithm presented in [3] and implemented in the promising tool `QCover`. We implemented our algorithm as a variant of `QCover` that we call `ICover`. Whereas `QCover` is based on invariants obtained from recent results on continuous Petri nets [9], our tool `ICover` is based on two classical methods: the *state equation* for Petri nets, and data-flow *sign analysis* [4]. On the 143 Petri net coverability questions that `QCover` solved, the tool `QCover` took 10318 seconds, while `ICover` used only 5517 seconds.

Outline. Section 2 recalls the Petri net coverability problem. Sections 3 and 4 present our backward coverability algorithm with pruning based on downward-closed invariants. In Sections 5 and 6, we recall two classical methods for computing invariants, namely the state (in-)equation and sign analysis. Section 7 is dedicated to the experimental evaluation of the tool `ICover`. In Section 8, we provide mathematical foundations for explaining our empirical good results based on the notion of limit-reachability in continuous Petri nets [18].

2 The Coverability Problem for Petri nets

A Petri net is a tuple $\mathcal{N} = (P, T, F, m_{init})$ comprising a finite set of *places* P, a finite set of *transitions* T disjoint of P, a *flow* function F from $(P \times T) \cup (T \times P)$ to $\mathbb{N}$, and an *initial* marking $m_{init} \in \mathbb{N}^P$. It is understood that $\mathbb{N}^P$ denotes the set of total maps from P to $\mathbb{N}$. Elements of $\mathbb{N}^P$ are called *markings*. Intuitively, a marking specifies how many *tokens* are in each place of the net. Tokens are consumed and produced through the firing of transitions. A transition $t \in T$ may fire only if it is enabled, meaning that each place p contains at least $F(p,t)$ tokens. Firing an enabled transition t modifies the contents of each place p by first removing $F(p,t)$ tokens and then adding $F(t,p)$ tokens. To clarify this intuitive description of the Petri net semantics, we introduce, for each transition $t \in T$, the t-step binary relation $\xrightarrow{t}$ over $\mathbb{N}^P$, defined by

$$m \xrightarrow{t} m' \;\Leftrightarrow\; \forall p \in P : m(p) \geq F(p,t) \wedge m'(p) = m(p) - F(p,t) + F(t,p)$$

The one-step binary relation $\rightarrow$ is the union of these t-step relations. Formally, $m \rightarrow m' \Leftrightarrow \exists t \in T : m \xrightarrow{t} m'$. The many-step binary relation $\xrightarrow{*}$ is the reflexive-transitive closure of $\rightarrow$.

Example 2.1. Figure 1 depicts a simple Petri net $\mathcal{N} = (P, T, F, m_{init})$ with places $P = \{p_1, p_2, p_3\}$, transitions $T = \{t_1, t_2, t_3\}$ and flow function F such

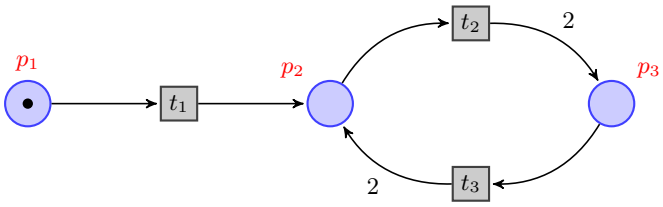

Fig. 1: Simple Petri net example

that $F(p_1, t_1) = 1$, $F(p_2, t_2) = 1$, $F(p_3, t_3) = 1$, $F(t_1, p_2) = 1$, $F(t_2, p_3) = 2$, $F(t_3, p_2) = 2$, and $F(p, t) = F(t, p) = 0$ for all other cases. The initial marking is $m_{init} = (1, 0, 0)$. The sequence of transitions $t_1 t_2 t_3$ may fire from the initial marking. Indeed, $(1, 0, 0) \xrightarrow{t_1} (0, 1, 0) \xrightarrow{t_2} (0, 0, 2) \xrightarrow{t_3} (0, 2, 1)$.

One of the most fundamental verification questions on Petri nets is coverability. In its simplest form, the coverability problem asks whether it is possible, by firing a sequence of transitions, to put a token in a given place. In essence, the coverability problem for Petri nets corresponds to the control-state reachability problem for other models of computation, such as counter machines, which are equipped with control states. The formal definition of coverability relies on a partial order over markings, defined hereafter.

Let $\leq$ denote the usual total order on $\mathbb{N}$. We extend $\leq$ over $\mathbb{N}^P$ componentwise, by $m \leq m' \Leftrightarrow \forall p \in P : m(p) \leq m'(p)$. This extension is a partial order over $\mathbb{N}^P$. Given two markings m and m' in $\mathbb{N}^P$, we say that m *covers* m' when $m \geq m'$. The *coverability problem* asks, given a Petri net $\mathcal{N} = (P, T, F, m_{init})$ and a *target* marking $m_{final} \in \mathbb{N}^P$, whether there exists a marking $m \in \mathbb{N}^P$ such that $m_{init} \xrightarrow{*} m$ and $m \geq m_{final}$. The main goal of this paper is to provide a simple, yet efficient procedure for solving this problem. Our method is inspired from the recent approach of [3]. Basically, we combine forward invariant generation techniques for Petri nets with backward reachability for well-structured transition systems [1,8]. Before delving into the details, we need some additional notations.

For a transition $t \in T$ and a set $S \subseteq \mathbb{N}^P$ of markings, we let $pre^t_{\mathcal{N}}(S)$ denote the predecessors of S via the transition t. Similarly, $pre_{\mathcal{N}}(S)$ and $pre^*_{\mathcal{N}}(S)$ denote the one-step and many-step predecessors of S, respectively. Formally, the functions $pre^t_{\mathcal{N}}$, $pre_{\mathcal{N}}$ and $pre^*_{\mathcal{N}}$ from $2^{\mathbb{N}^P}$ to $2^{\mathbb{N}^P}$ are defined by

$$\begin{aligned} pre^t_{\mathcal{N}}(S) &= \{m \in \mathbb{N}^P \mid \exists m' \in S : m \xrightarrow{t} m'\} \\ pre_{\mathcal{N}}(S) &= \{m \in \mathbb{N}^P \mid \exists m' \in S : m \to m'\} \\ pre^*_{\mathcal{N}}(S) &= \{m \in \mathbb{N}^P \mid \exists m' \in S : m \xrightarrow{*} m'\} \end{aligned}$$

Given a subset $S \subseteq \mathbb{N}^P$ of markings, we let $\uparrow S$ and $\downarrow S$ denote its *upward closure* and *downward closure*, respectively. These are defined by

$$\begin{aligned} \uparrow S &= \{u \in \mathbb{N}^P \mid \exists m \in S : u \geq m\} \\ \downarrow S &= \{d \in \mathbb{N}^P \mid \exists m \in S : d \leq m\} \end{aligned}$$

A subset $S \subseteq \mathbb{N}^P$ is called *upward-closed* when $S = \uparrow S$, and it is called *downward-closed* when $S = \downarrow S$.

Notation 2.2. For the remainder of the paper, to avoid clutter, we will simply write m in place of $\{m\}$ for singletons, when this causes no confusion.

Recall that the coverability problem asks whether $m_{init} \xrightarrow{*} m \geq m_{final}$ for some marking $m \in \mathbb{N}^P$. This problem is equivalently phrased as the question whether m_{init} belongs to $pre^*_{\mathcal{N}}(\uparrow m_{final})$. This formulation can be seen as a backward analysis question. We may also phrase the coverability problem in terms of a forward analysis question, using the notion of coverability set.

Given a Petri net $\mathcal{N} = (P, T, F, m_{init})$, the *coverability set* of $\mathcal{N}$ is the set $Cov_{\mathcal{N}} = \downarrow\{m \in \mathbb{N}^P \mid m_{init} \xrightarrow{*} m\}$. It is readily seen that the coverability problem is equivalent to the question whether m_{final} belongs to $Cov_{\mathcal{N}}$. We are now equipped with the necessary notions to present our mixed forward/backward approach for the coverability problem.

3 Backward Coverability Analysis with Pruning

We now present our method to solve the coverability problem for Petri nets. This section gives the mathematical foundations of our approach, with no regard for implementability. We will focus on the implementation of this approach in Section 4.

The classical backward reachability approach for the coverability problem [1,8] consists in computing a growing sequence $U_0 \subseteq U_1 \subseteq \cdots$ of upward-closed subsets of $\mathbb{N}^P$ that converges to $pre^*_{\mathcal{N}}(\uparrow m_{final})$. Here, we modify this growing sequence in order to leverage an a priori known over-approximation of the coverability set. In practice, this means that we narrow the backward reachability search by pruning some markings that are known to be not coverable.

An *invariant* for a Petri net $\mathcal{N} = (P, T, F, m_{init})$ is any subset $I \subseteq \mathbb{N}^P$ that contains every reachable marking, i.e., every marking m with $m_{init} \xrightarrow{*} m$. Observe that a downward-closed subset of $\mathbb{N}^P$ is an invariant of $\mathcal{N}$ if, and only if, it contains $Cov_{\mathcal{N}}$. Sections 5 and 6 will discuss the automatic generation of downward-closed invariants.

For the remainder of this section, we consider a Petri net $\mathcal{N} = (P, T, F, m_{init})$ and we assume that we are given a downward-closed invariant I for $\mathcal{N}$. We introduce the sequence $U_0, U_1, \ldots$ of subsets of $\mathbb{N}^P$ defined as follows:

$$\begin{aligned} U_0 &= \uparrow(m_{final} \cap I) \\ U_{k+1} &= \uparrow(pre_{\mathcal{N}}(U_k) \cap I) \cup U_k \end{aligned}$$

Observe that each U_k is upward-closed and that the sequence $(U_k)_k$ is growing for inclusion. On the contrary to the classical backward reachability approach [1,8], U_{k+1} does not consider all one-step predecessors of U_k, but discards those that are not in I. Note that by taking $I = \mathbb{N}^P$, which is trivially a downward-closed

invariant, we obtain the same growing sequence as in the classical backward reachability approach [1,8]. The two following lemmas show that we can use the sequence $(U_k)_k$ to solve the coverability problem.

Lemma 3.1. *The sequence $(U_k)_k$ is ultimately stationary.*

Lemma 3.2. *It holds that $m_{final} \in Cov_{\mathcal{N}}$ if, and only if, $m_{init} \in \bigcup_k U_k$.*

We have presented in this section a growing sequence of upward-closed subsets of markings that is ultimately stationary and whose limit contains enough information to solve the coverability problem. Our next step is to transform this sequence into an algorithm.

4 The `ICover` Algorithm

In this section, we turn the growing sequence $(U_k)_k$ of upward-closed subsets of markings defined in Section 3 into an algorithm. Of course, we cannot directly compute the sets U_k since they may be infinite (in fact, they are either empty or infinite). Instead, we will compute finite sets $B_k \subseteq \mathbb{N}^P$ such that $U_k = {\uparrow}B_k$. The existence of such finite sets is guaranteed by the following lemma. A *basis* of an upward-closed subset $U \subseteq \mathbb{N}^P$ is any set $B \subseteq \mathbb{N}^P$ such that $U = {\uparrow}B$. Recall that a *minimal* element of a subset $S \subseteq \mathbb{N}^P$ is any $m \in S$ such that $u \leq m \Rightarrow u = m$ for every $u \in S$.

Lemma 4.1. *For every subset $S \subseteq \mathbb{N}^P$, the set $\operatorname{Min} S$ of its minimal elements is finite and satisfies ${\uparrow}S = {\uparrow}\operatorname{Min} S$.*

Corollary 4.2. *Every upward-closed subset $U \subseteq \mathbb{N}^P$ admits a finite basis.*

We still need to show how to compute a finite basis of U_{k+1} from a finite basis of U_k. To this end we introduce, for each transition $t \in T$, the *covering predecessor* function $cpre^t_{\mathcal{N}} : \mathbb{N}^P \to \mathbb{N}^P$ defined by

$$cpre^t_{\mathcal{N}}(m)(p) \;=\; F(p,t) + \max(0, m(p) - F(t,p))$$

Informally, $cpre^t_{\mathcal{N}}(m)$ is the least marking that can cover m in one step by firing the transition t. This property will be formally stated in Lemma 4.3. The function $cpre^t_{\mathcal{N}}$ is extended to sets of markings by $cpre^t_{\mathcal{N}}(S) = \{cpre^t_{\mathcal{N}}(m) \mid m \in S\}$.

Lemma 4.3. *It holds that $pre^t_{\mathcal{N}}({\uparrow}m) = {\uparrow}cpre^t_{\mathcal{N}}(m)$ for every marking $m \in \mathbb{N}^P$.*

The previous lemma can easily be extended to sets of markings. We extend it further, in Lemma 4.4, to bridge the gap with the definition of $(U_k)_k$. The lemma shows how to compute a finite basis of U_{k+1} from a finite basis of U_k.

Lemma 4.4. *Let I be a downward-closed invariant for $\mathcal{N}$. For every subset $S \subseteq \mathbb{N}^P$, it holds that ${\uparrow}pre^t_{\mathcal{N}}(({\uparrow}S) \cap I) = {\uparrow}(cpre^t_{\mathcal{N}}(S) \cap I)$.*

`ICover`$(\mathcal{N}, m_{final}, I)$

Input: A Petri Net $\mathcal{N} = (P, T, F, m_{init})$, a target marking $m_{final} \in \mathbb{N}^P$ and a downward-closed invariant I for $\mathcal{N}$.

Output: Whether there exists a marking $m \in \mathbb{N}^P$ such that $m_{init} \xrightarrow{*} m$ and $m \geq m_{final}$.

```
1  begin
2      if m_final ∈ I then
3          B ← {m_final}
4      else
5          B ← ∅
6      while m_init ∉ ↑B do
7          N ← {cpre^t_N(m) | t ∈ T, m ∈ B} \ ↑B        /* new predecessors */
8          P ← N ∩ I                                    /* prune uncoverable markings */
9          if P = ∅ then
10             return False
11         B ← Min(B ∪ P)
12     return True
```

Proof. The straightforward extension of Lemma 4.3 to sets of markings shows that $pre^t_{\mathcal{N}}(\uparrow S) = \uparrow cpre^t_{\mathcal{N}}(S)$ for every subset $S \subseteq \mathbb{N}^P$. Moreover, it is readily seen that, for every subset $S \subseteq \mathbb{N}^P$, $\uparrow((\uparrow S) \cap I) = \uparrow(S \cap I)$. This property follows from the assumption that I is downward-closed. We derive that

$$\begin{aligned} \uparrow(cpre^t_{\mathcal{N}}(S) \cap I) &= \uparrow((\uparrow cpre^t_{\mathcal{N}}(S)) \cap I) \\ &= \uparrow(pre^t_{\mathcal{N}}(\uparrow S) \cap I) \end{aligned}$$

This concludes the proof of the lemma. □

The previous lemma leads to a backward coverability algorithm, called `ICover` and presented on page 6. Basically, this procedure symbolically computes the growing sequence $(U_k)_k$ of upward-closed sets. Let us make the relationship between the procedure and the sequence $(U_k)_k$ more precise. Consider an input instance $(\mathcal{N}, m_{final}, I)$ of `ICover`. Since the procedure is deterministic, `ICover`$(\mathcal{N}, m_{final}, I)$ has a unique maximal execution, that either terminates (at line 10 or 12) or iterates the **while** loop (lines 6–11) indefinitely. Let $\ell_B, \ell_P \in \mathbb{N} \cup \{\infty\}$ denote the numbers of executions of lines 6 and 9, respectively. It is understood that $l_P \leq l_B \leq l_P + 1$, with the convention that $\infty + 1 = \infty$. Let $(B_k)_{k<\ell_B}$ and $(P_k)_{k<\ell_P}$ denote the successive values at lines 6 and 9 of the variables B and P, respectively.

Lemma 4.5. *For every k with $0 \leq k < \ell_B$, the set B_k is a finite basis of U_k. For every k with $0 \leq k < \ell_P$, the set P_k is a finite basis of $\uparrow(U_{k+1} \setminus U_k)$.*

Theorem 4.6. *The procedure* `ICover` *terminates on every input and is correct.*

Remark 4.7. Petri nets obtained by translation from high-level concurrent programs often contain transitions that cannot be fired from any reachable marking. Downward-closed invariants can be used in a pre-processing algorithm to filter out some of them. Basically, if a transition t is not enabled in any marking of an invariant I, it can be safely removed without modifying the coverability set. Algorithmically, when I is downward-closed, detecting such a property just reduces to a membership problem in I. In fact a transition t is enabled in a downward-closed set of markings D if, and only if, D contains the marking m_t defined by $m_t(p) = F(p,t)$ for every place p.

The algorithm `ICover` is parametrized by an a priori known downward-closed invariant that is given as input. On the one hand, this invariant needs to be precise enough to discard markings (at line 8) and accelerate the main loop. On the other hand, we need to decide efficiently whether a marking is in the invariant, to avoid slowing down the main loop. The next two sections show how to generate downward-closed invariants with efficient membership testing.

5 State Inequation for Downward-Closed Invariants

The state equation provides a simple over-approximation of Petri net reachability relations that was successfully used in two recent algorithms for deciding the coverability problems [7,3]. This equation is obtained by introducing the total function $\Delta(t)$ in $\mathbb{Z}^P$ called the *displacement* of a transition t and defined for every place p by $\Delta(t)(p) = F(t,p) - F(p,t)$. Let us assume that a marking m_{final} is in the coverability set of a Petri net $\mathcal{N}$. It follows that there exists a word $t_1 \dots t_k$ of transitions and a marking $m \geq m_{final}$ such that $m_{init} \xrightarrow{t_1} \cdots \xrightarrow{t_k} m$. We derive the following relation:

$$m_{init} + \Delta(t_1) + \cdots + \Delta(t_k) = m \geq m_{final}$$

By reordering the sum $\Delta(t_1) + \cdots + \Delta(t_k)$, we can group together the displacements $\Delta(t)$ corresponding to the same transition t. Denoting by $\lambda(t)$ the number of occurrences of t in the word $t_1 \dots t_k$, we get:

$$m_{init} + \sum_{t \in T} \lambda(t)\Delta(t) \geq m_{final} \tag{1}$$

The relation (1) is called the *state inequation* for the coverability problem. Notice that a similar equation can be derived for the reachability problem by replacing the inequality by an equality. We do not consider this equality in the sequel since we restrict our attention to the coverability problem. We introduce the following set I_S where $\mathbb{Q}_{\geq 0}$ is the set of non-negative rational numbers.

$$I_S = \{m \in \mathbb{N}^P \mid \exists \lambda \in \mathbb{Q}_{\geq 0}^T : m_{init} + \sum_{t \in T} \lambda(t)\Delta(t) \geq m\} \tag{2}$$

Proposition 5.1. *The set I_S is a downward-closed invariant with a polynomial-time membership problem.*

A more precise downward-closed invariant can be obtained by requiring that $\lambda \in \mathbb{N}^T$. In particular, the pruned backward algorithm presented in Section 4 should produce smaller sets of configurations with this more precise invariant. In practice, we do not observe any significant improvement on a large set of benchmarks. Moreover, whereas the membership problem of a marking m is decidable in polynomial time when λ ranges over $\mathbb{Q}^T_{\geq 0}$, the problem becomes NP-complete when λ is restricted to $\mathbb{N}^T$.

6 Sign Analysis for Downward-Closed Invariants

In this section we introduce a downward-closed invariant based on data-flow sign analysis [4]. Rephrased in the context of Petri nets, an invariant I is said to be *inductive* if $m \xrightarrow{t} m'$ and $m \in I$ implies $m' \in I$. Sign analysis then reduces to the computation of the maximal (for the inclusion) set Z of places such that the following set I_Z is an inductive invariant:

$$I_Z = \{m \in \mathbb{N}^P \mid \bigwedge_{p \in Z} m(p) = 0\} \tag{3}$$

The unicity of that set is immediate since the class of sets Z such that I_Z is an invariant is clearly closed under union. In the sequel, Z denotes the maximal set satisfying this property, and this maximal set is shown to be computable in polynomial time thanks to a fixpoint propagation. We introduce the operator $\mathrm{prop}_t : 2^P \to 2^P$ associated to a transition t and defined for any set Q of places as follows:

$$\mathrm{prop}_t(Q) = \begin{cases} \{q \in P \mid F(t,q) > 0\} & \text{if } \bigwedge_{p \in P \setminus Q} F(p,t) = 0 \\ \emptyset & \text{otherwise} \end{cases}$$

Intuitively, if t is a transition such that $\bigwedge_{p \in P \setminus Q} F(p,t) = 0$ then from a marking with large number of tokens in each place of Q, it is possible to fire t. In particular places q satisfying $F(t,q) > 0$ cannot be in Z. This property is formally stated by the following lemma.

Lemma 6.1. *We have* $\mathrm{prop}_t(Q) \subseteq P \setminus Z$ *for every set* $Q \subseteq P \setminus Z$.

The set Z can be computed as a fixpoint by introducing the non-decreasing sequence $Q_0, Q_1, \ldots$ of places defined as follows:

$$Q_0 = \{q \in P \mid m_{init}(q) > 0\}$$
$$Q_{k+1} = Q_k \cup \bigcup_{t \in T} \mathrm{prop}_t(Q_k)$$

Let us notice that the set $Q = \bigcup_{k \geq 0} Q_k$ is computable in polynomial time. The following lemma shows that Q provides the set Z as a complement.

Lemma 6.2. *We have* $Z = P \setminus Q$.

Corollary 6.3. *The set* Z *is computable in polynomial time.*

7 Experimental Evaluation

We implemented our approach using the `QCover` [3] tool as a starting point. This tool, which implements a backward coverability algorithm for Petri nets, is written in Python and relies on the SMT-solver $z3$ [16]. `QCover` also uses some other heuristics that we kept unchanged. `QCover` was competitive with others tools especially for uncoverable Petri net. Only the `BFC` tool performs significantly better on coverable Petri net. We have made two modifications to `QCover`. First, we have added a pre-processing step (see Remark 4.7) based on sign analysis. Second, we have replaced their pruning technique, which is based on coverability in continuous Petri nets, by the one of our algorithm `ICover` presented in Section 4. `ICover` is available as a patch [11] for `QCover` [2].

To test our implementation, we used the same benchmark as `Petrinizer` [7] and `QCover` [3]. It comprises models from various sources: `Mist` [10], `BFC` [13], `Erlang` programs abstracted into Petri nets [6], as well as so-called `medical` and `bug_tracking` examples [7]. We let each tool work for 2000 seconds in a machine on Ubuntu Linux 14.04 with Intel(R) Core(TM) i7-4770 CPU at 3.40GH with 16 GB of memory for each benchmark. The computation times are the sum of the `system` and `user` times. Overall `QCover` solved 106 uncoverable instances on 115 Petri net and 37 coverable problems on 61 Petri nets. `ICover` was able to find one more coverable instance. In fact calling `QCover` on the Petri net computed by the pre-processing, that we will call `QCover/Pp`, can even solve one more uncoverable instance than `ICover`. On the 143 instances that `QCover` solved, the tool took 10318 seconds, `QCover/Pp` used 6479 seconds, and `ICover` used only 5162 seconds.

Figure 2(a) shows the comparison between `ICover` and `QCover` in time. The straight line represents when the two tools took the same time. Each dot represents a coverability question. When the dot is under the line, it means that `ICover` was faster than `QCover` and conversely. There are three instances where `QCover` performs very well, under a second, and where `ICover` took a few tens of seconds to answer. For the three cases, the formula used by `QCover` for coverability in $\mathbb{Q}$ was enough to discard the target as uncoverable and it didn't have to enter in the while loop. But `ICover` wasn't able to discard the target and had to enter the while loop in the three cases. We also see two dots above the line at the middle of the figure. The pre-processing took respectively 12 and 45 seconds while the initial Petri net was solved by `QCover` in respectively 16 and 33 seconds. The pre-processing has not been optimized yet, and it could probably run faster.

Figure 2(b) and (c) show the intermediate comparisons: `ICover` versus `QCover/Pp` and `QCover/Pp` versus `QCover`. We can observe that the pre-processing has a major impact on the good performance of `ICover` compared to `QCover`.

Figure 2(e) and Figure 2(f) aims to show the effect of the pre-processing on the size of Petri nets. The former show the percentage of places left after pre-processing. Some Petri nets kept all their places but others were left with only 2.5% of their initial places. And most of Petri nets lost a significant

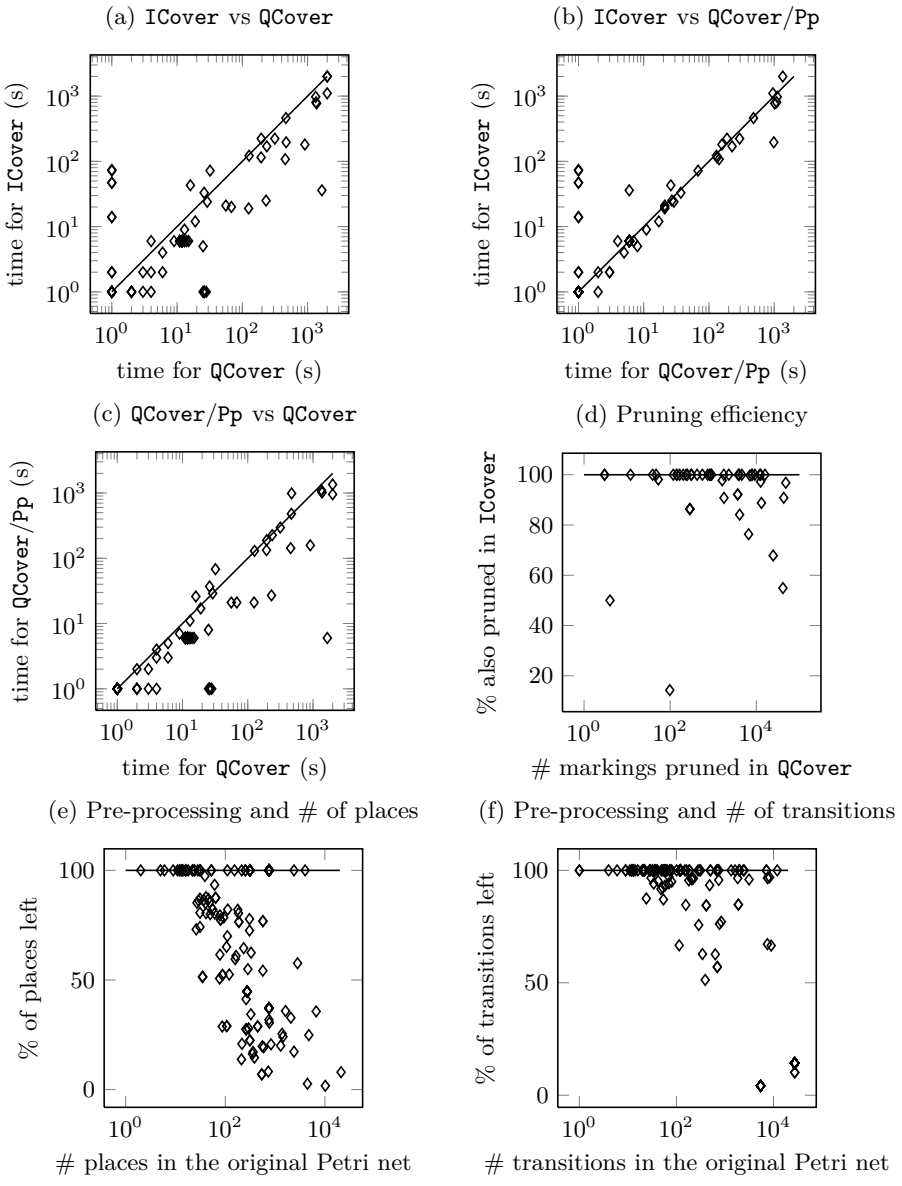
(a) ICover vs QCover
time for ICover (s)
time for QCover (s)
(b) ICover vs QCover/Pp
time for ICover (s)
time for QCover/Pp (s)
(c) QCover/Pp vs QCover
time for QCover/Pp (s)
time for QCover (s)
(d) Pruning efficiency
% also pruned in ICover
# markings pruned in QCover
(e) Pre-processing and # of places
% of places left
# places in the original Petri net
(f) Pre-processing and # of transitions
% of transitions left
# transitions in the original Petri net

Fig. 2: Experimental results for ICover, QCover and QCover/Pp

number of places. The latter shows the percentages of transitions left after the pre-processing. Overall less transitions were cut than places. Half of the Petri nets kept all their transitions, but some were left with only 4% of their initial transitions.

Figure 2(d) compares the efficiency of pruning between `ICover` and `QCover`. Again, each dot represents a coverability question. As discussed in Section 8, `QCover` always prunes at least as many markings as `ICover` (but at the expense of more complex pruning tests). A value of 100% means that `ICover` was able to prune the same markings as `QCover`. It turns out that on most instances, this perfect value of 100% is obtained. This is rather surprising at first sight, and warrants an investigation, which is the focus of the next section.

8 Comparison with Continuous Petri Net

Continuous Petri nets are defined like Petri nets except that transitions can be fired a non-negative rational number of times. The firing of such a transition produces markings with non-negative rational numbers of tokens. Under such a semantics, called the *continuous semantics*, the reachability problem was recently proved to be decidable in polynomial time [9]. Based on this observation, the tool `QCover` implements the pruning backward coverability algorithm presented in Section 3 with a downward-closed invariant derived from the continuous semantics. Whereas this invariant is more precise than the downward-closed invariant obtained from the state inequation introduced in Section 5, we have seen in Section 7 that such an improvement is overall not useful in practice for the pruning backward algorithm. In this section, we provide a simple structural condition on Petri nets in such a way the two kinds of downward-closed invariants derived respectively from the continuous semantics and the state inequation are "almost" equal. This structural condition is shown to be natural since it is fulfilled by the Petri nets obtained after the pre-processing introduced in Remark 4.7.

A *continuous marking* is a mapping $m \in \mathbb{Q}_{\geq 0}^P$ where $\mathbb{Q}_{\geq 0}$ denotes the set of non-negative rational numbers, and P the set of places. Given $r \in \mathbb{Q}_{\geq 0}$ and a transition t, the continuous rt-step binary relation $\xrightarrow{rt}$ over the continuous markings is defined by

$$m \xrightarrow{rt} m' \;\Leftrightarrow\; \forall p \in P : m(p) \geq r.F(p,t) \wedge m'(p) = m(p) - r.F(p,t) + r.F(t,p)$$

The one-step continuous binary relation $\dashrightarrow$ is the union of these rt-step relations. Formally, $m \dashrightarrow m'$ if there exists $r \in \mathbb{Q}_{\geq 0}$ and $t \in T$ such that $m \xrightarrow{rt} m'$. The many-step continuous binary relation $\xrightarrow{*}$ is the reflexive-transitive closure of $\dashrightarrow$. We also introduce the binary relation $\xrightarrow{\infty}$ defined over the continuous markings by $m \xrightarrow{\infty} m'$ if there exists a sequence $(m_k)_{k\geq 0}$ of continuous markings that *converges* towards m' with the classical topology on $\mathbb{Q}_{\geq 0}^P$ and such that $m \xrightarrow{*} m_k$ for every k.

Example 8.1. Let us look back at the simple Petri net $\mathcal{N}$ depicted in Figure 1. For every positive natural number k, we have:

$$(1,0,0) \xrightarrow{\frac{1}{k}t_1} (1-\frac{1}{k},\frac{1}{k},0) \xrightarrow{\frac{1}{k}t_2\frac{1}{k}t_3} (1-\frac{1}{k},\frac{2}{k},\frac{1}{k}) \cdots \xrightarrow{\frac{1}{k}t_2\frac{1}{k}t_3} (1-\frac{1}{k},1+\frac{1}{k},1)$$

It follows that $(1,0,0) \overset{\infty}{\dashrightarrow} (1,1,1)$. Notice that the relation $(1,0,0) \overset{*}{\dashrightarrow} (1,1,1)$ does not hold.

The downward-closed invariant used in the tool `QCover` for implementing the pruning backward algorithm is defined as follows:

$$I_C = \{m \in \mathbb{N}^P \mid \exists m' \in \mathbb{Q}^P_{\geq 0} : m_{init} \overset{*}{\dashrightarrow} m' \geq m\} \tag{4}$$

Recall that in Section 5 we introduced the set I_S for denoting the downward-closed invariant derived from the state inequation. The following result[1] provides a characterization of that invariant when the Petri net satisfies a structural condition.

Theorem 8.2 ([18, Theorem 7]). *If every transition is fireable from the downward-closed invariant I_Z introduced in Section 6, we have:*

$$I_S = \{m \in \mathbb{N}^P \mid \exists m' \in \mathbb{Q}^P_{\geq 0} : m_{init} \overset{\infty}{\dashrightarrow} m' \geq m\} \tag{5}$$

The two equalities Equation (4) and Equation (5) show that I_S and I_C are very similar for Petri nets satisfying the structural condition stated in Theorem 8.2. This condition will be fulfilled by the Petri nets produced by the pre-processing algorithm introduced in Remark 4.7. Notice that even if the membership problem in I_S and I_C are both decidable in polynomial time, the extra computational cost for deciding the membership problem for the invariant I_C, even for efficient SMT solvers like `Z3`, is not neglectable. Naturally, if a marking is in I_C then it is also in I_S, and the converse property is false in general as shown by Example 8.1. However, in practice, we observed that configurations that are in I_S are very often also in I_C (see Figure 2(d)), as already mentioned in Section 7.

9 Conclusion

Petri nets have recently been used as low-level models for model-checking concurrent systems written in high-level programming languages [5,6]. The original verification question on the concurrent program reduces to a coverability question on the resulting Petri net. We have proposed in this paper a family of simple coverability algorithms parametrized by downward-closed invariants. As future work, we intend to look for classes of downward-closed invariants with a good tradeoff between precision and efficient membership.

References

1. Abdulla, P.A., Cerans, K., Jonsson, B., Tsay, Y.: Algorithmic analysis of programs with well quasi-ordered domains. Information and Computation 160(1-2), 109–127 (2000)

[1] The statement of Theorem 7 in [18] is wrong since it is based on a too strong definition of limit-reachability. However, the proof becomes correct with our definitions and notations.

2. Blondin, M., Finkel, A., Haase, C., Haddad, S.: QCover with benchmarks. http://www-etud.iro.umontreal.ca/~blondimi/doc/qcover_with_benchmarks.zip
3. Blondin, M., Finkel, A., Haase, C., Haddad, S.: Approaching the coverability problem continuously. In: TACAS. pp. 480–496. Springer (2016)
4. Cousot, P., Cousot, R.: Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: POPL. pp. 238–252. ACM (1977)
5. Donaldson, A., Kaiser, A., Kroening, D., Wahl, T.: Symmetry-aware predicate abstraction for shared-variable concurrent programs. In: CAV. pp. 356–371. Springer (2011)
6. D'Osualdo, E., Kochems, J., Ong, C.L.: Automatic verification of Erlang-style concurrency. In: SAS. pp. 454–476. Springer (2013)
7. Esparza, J., Ledesma-Garza, R., Majumdar, R., Meyer, P.J., Niksic, F.: An SMT-based approach to coverability analysis. In: CAV. pp. 603–619. Springer (2014)
8. Finkel, A., Schnoebelen, P.: Well-structured transition systems everywhere! Theoretical Computer Science 256(1-2), 63–92 (2001)
9. Fraca, E., Haddad, S.: Complexity analysis of continuous Petri nets. Fundamenta Informaticae 137(1), 1–28 (2015)
10. Ganty, P.: Mist – A safety checker for petri nets and extensions. http://github.com/pierreganty/mist
11. Geffroy, T., Leroux, J., Sutre, G.: ICover patch. http://dept-info.labri.u-bordeaux.fr/~tgeffroy/icover/
12. German, S.M., Sistla, A.P.: Reasoning about systems with many processes. Journal of the Association for Computing Machinery 39(3), 675–735 (1992)
13. Kaiser, A., Kroening, D., Wahl, T.: A widening approach to multithreaded program verification. ACM Trans. Program. Lang. Syst. 36(4), 14:1–14:29 (2014)
14. Karp, R.M., Miller, R.E.: Parallel program schemata. Journal of Computer and System Sciences 3(2), 147–195 (1969)
15. Lipton, R.J.: The reachability problem requires exponential space. Tech. Rep. 62, Yale University (1976)
16. de Moura, L.M., Bjørner, N.: Z3: an efficient SMT solver. In: TACAS. pp. 337–340. Springer (2008)
17. Rackoff, C.: The covering and boundedness problems for vector addition systems. Theoretical Computer Science 6(2), 223–231 (1978)
18. Recalde, L., Teruel, E., Silva, M.: Autonomous continuous P/T systems. In: ICATPN. pp. 107–126. Springer (1999)

A Proofs for Section 3

Lemma 3.1. *The sequence $(U_k)_k$ is ultimately stationary.*

Proof. The partial order $\leq$ on $\mathbb{N}^P$ is a well-quasi-order by Dickson's Lemma. Therefore, every growing sequence of upward-closed sets is ultimately stationary (see, e.g., [8, Lemma 2.4]). The lemma follows from the observation that each U_k is upward-closed and that the sequence $(U_k)_k$ is growing. □

Lemma 3.2. *It holds that $m_{final} \in Cov_{\mathcal{N}}$ if, and only if, $m_{init} \in \bigcup_k U_k$.*

Proof. If $m_{final} \in Cov_{\mathcal{N}}$ then $m_{init} \xrightarrow{*} m \geq m_{final}$ for some marking m in $\mathbb{N}^P$. Since $m_{init} \xrightarrow{*} m$, there exists $m_0, \ldots, m_n \in \mathbb{N}^P$ such that $m_{init} = m_n$, $m_n \to m_{n-1} \cdots \to m_0$ and $m_0 \geq m_{final}$. First observe that $m_i \in I$ for every $i \in \{0, \ldots, m\}$ because I is an invariant for $\mathcal{N}$. Moreover, $m_{final} \in I$ since I is downward-closed, $m_0 \geq m_{final}$ and $m_0 \in I$. We prove, by induction on i, that $m_i \in U_i$ for all $i \in \{0, \ldots, n\}$. The basis $m_0 \in U_0$ follows from the facts that $m_0 \geq m_{final}$ and $m_{final} \in I$. For the induction step, let $i \in \{0, \ldots, n-1\}$ and assume that $m_i \in U_i$. Recall that $m_{i+1} \in I$ and $m_{i+1} \to m_i$. It follows that $m_{i+1} \in (pre_{\mathcal{N}}(U_i) \cap I) \subseteq U_{i+1}$. We have thus shown that $m_n \in U_n$, hence, $m_{init} = m_n$ belongs to $\bigcup_k U_k$.

Let us show the converse of the lemma. We first prove, by induction on k, that $U_k \subseteq pre^*_{\mathcal{N}}(\uparrow m_{final})$ for every $k \in \mathbb{N}$. The basis follows from the observation that $U_0 \subseteq \uparrow m_{final} \subseteq pre^*_{\mathcal{N}}(\uparrow m_{final})$. For the induction step, let $k \in \mathbb{N}$ and assume that $U_k \subseteq pre^*_{\mathcal{N}}(\uparrow m_{final})$. Recall that $U_{k+1} = \uparrow(pre_{\mathcal{N}}(U_k) \cap I) \cup U_k$, hence, $U_{k+1} \subseteq \uparrow pre_{\mathcal{N}}(U_k) \cup U_k$. Since $\mathcal{N}$ is a Petri net, $pre_{\mathcal{N}}(S)$ is upward-closed for every upward-closed subset $S \subseteq \mathbb{N}^P$. It follows that $\uparrow pre_{\mathcal{N}}(U_k) = pre_{\mathcal{N}}(U_k)$. We derive from the induction hypothesis that $U_{k+1} \subseteq pre_{\mathcal{N}}(U_k) \cup U_k \subseteq pre^*_{\mathcal{N}}(\uparrow m_{final})$. We have thus shown that $U_k \subseteq pre^*_{\mathcal{N}}(\uparrow m_{final})$ for every $k \in \mathbb{N}$. The observation that $m_{final} \in Cov_{\mathcal{N}} \Leftrightarrow m_{init} \in pre^*_{\mathcal{N}}(\uparrow m_{final})$ concludes the proof of the lemma. □

B Proofs for Section 4

Lemma 4.1. *For every subset $S \subseteq \mathbb{N}^P$, the set $\operatorname{Min} S$ of its minimal elements is finite and satisfies $\uparrow S = \uparrow \operatorname{Min} S$.*

Proof. The partial order $\leq$ on $\mathbb{N}^P$ is a well-quasi-order by Dickson's Lemma. Therefore, the set $\operatorname{Min} S$ of minimal elements of S is necessarily finite. Moreover, $S \subseteq \uparrow \operatorname{Min} S$ since $\leq$ is well-founded. It follows that $\uparrow S = \uparrow \operatorname{Min} S$. □

Lemma 4.3. *It holds that $pre^t_{\mathcal{N}}(\uparrow m) = \uparrow cpre^t_{\mathcal{N}}(m)$ for every marking $m \in \mathbb{N}^P$.*

Proof. Let $u \in pre^t_{\mathcal{N}}(\uparrow m)$. There exists $v \geq m$ such that $u \xrightarrow{t} v$. Consider a place $p \in P$. It holds that $u(p) \geq F(p,t)$ and $v(p) = u(p) - F(p,t) + F(t,p)$ since $u \xrightarrow{t} v$. We consider two cases.

– If $m(p) \leq F(p,t)$ then $cpre^t_{\mathcal{N}}(m)(p) = F(p,t) \leq u(p)$.
– If $m(p) \geq F(p,t)$ then $cpre^t_{\mathcal{N}}(m)(p) = F(p,t) + m(p) - F(t,p)$. Since $v \geq m$, we get that $cpre^t_{\mathcal{N}}(m)(p) \leq F(p,t) + v(p) - F(t,p) = u(p)$.

In both cases, we obtain that $cpre^t_{\mathcal{N}}(m)(p) \leq u(p)$. We have thus shown that $u \in \uparrow cpre^t_{\mathcal{N}}(m)$.

Conversely, let $u \in \uparrow cpre^t_{\mathcal{N}}(m)$. This means that $u(p) \geq cpre^t_{\mathcal{N}}(m)(p)$ for every place $p \in P$. Therefore, $u(p) \geq F(p,t)$ and $u(p) \geq F(p,t) + m(p) - F(t,p)$. It follows that $u \xrightarrow{t} v$ for the marking $v \geq m$ defined by $v(p) = u(p) - F(p,t) + F(t,p)$. We have thus shown that $u \in pre^t_{\mathcal{N}}(\uparrow m)$. □

Lemma 4.5. *For every k with $0 \leq k < \ell_B$, the set B_k is a finite basis of U_k. For every k with $0 \leq k < \ell_P$, the set P_k is a finite basis of $\uparrow(U_{k+1} \setminus U_k)$.*

Proof. It is readily seen that B_k and P_k are finite subsets of $\mathbb{N}^P$ for every k. We first observe that, for every k with $0 \leq k < \ell_P$,

$$\begin{array}{rcll} \uparrow P_k & = & \uparrow((\{cpre^t_{\mathcal{N}}(m) \mid t \in T, m \in B_k\} \setminus \uparrow B_k) \cap I) & \text{[Lines 7–8]} \\ & = & \uparrow(\{cpre^t_{\mathcal{N}}(m) \mid t \in T, m \in B_k\} \cap (I \setminus \uparrow B_k)) & \\ & = & \bigcup_{t \in T} \uparrow(cpre^t_{\mathcal{N}}(B_k) \cap (I \setminus \uparrow B_k)) & \\ & = & \bigcup_{t \in T} \uparrow(pre^t_{\mathcal{N}}(\uparrow B_k) \cap (I \setminus \uparrow B_k)) & \text{[Lemma 4.4]} \\ & = & \uparrow(pre_{\mathcal{N}}(\uparrow B_k) \cap (I \setminus \uparrow B_k)) & \\ & = & \uparrow((pre_{\mathcal{N}}(\uparrow B_k) \cap I) \setminus \uparrow B_k) & \end{array}$$

Let us now prove, by induction on k, that $U_k = \uparrow B_k$ for every k with $0 \leq k < \ell_B$. The basis $U_0 = \uparrow B_0$ follows from lines 1–6 of `ICover` and from the definition of U_0. For the induction step, let $k \in \mathbb{N}$ with $k+1 < \ell_B$, and assume that $U_k = \uparrow B_k$. Line 11 entails that $B_{k+1} = \mathrm{Min}(B_k \cup P_k)$. It follows that

$$\begin{array}{rcll} \uparrow B_{k+1} & = & \uparrow B_k \cup \uparrow P_k & \text{[Lemma 4.1]} \\ & = & \uparrow B_k \cup \uparrow((pre_{\mathcal{N}}(\uparrow B_k) \cap I) \setminus \uparrow B_k) & \\ & = & U_k \cup \uparrow((pre_{\mathcal{N}}(U_k) \cap I) \setminus U_k) & [U_k = \uparrow B_k] \\ & = & U_k \cup \uparrow(pre_{\mathcal{N}}(U_k) \cap I) & \\ & = & U_{k+1} & \end{array}$$

This concludes the proof that $U_k = \uparrow B_k$ for every k with $0 \leq k < \ell_B$. Moreover, coming back to the characterization of $\uparrow P_k$, we get that

$$\begin{array}{rcl} \uparrow P_k & = & \uparrow((pre_{\mathcal{N}}(\uparrow B_k) \cap I) \setminus \uparrow B_k) \\ & = & \uparrow((pre_{\mathcal{N}}(U_k) \cap I) \setminus U_k) \\ & = & \uparrow(U_{k+1} \setminus U_k) \end{array}$$

for every k with $0 \leq k < \ell_P$. □

Theorem 4.6. *The procedure* `ICover` *terminates on every input and is correct.*

Proof. Let us first prove termination. We need to show that the unique maximal execution of $\texttt{ICover}(\mathcal{N}, m_{final}, I)$ is finite. By contradition, assume that $\ell_B = \infty$.

According to Lemma 3.1, there exists an index $h \in \mathbb{N}$ such that $U_h = U_{h+1}$. We derive from Lemma 4.5 that $P_h = \emptyset$. Therefore, the execution should terminate at line 10 during the $(h+1)^{\text{th}}$ iteration of the **while** loop. This contradicts our assumption that $\ell_B = \infty$.

We now turn our attention to the correctness of `ICover`. As it is finite, the unique maximal execution of $\texttt{ICover}(\mathcal{N}, m_{final}, I)$ either returns `False` at line 10 or returns `True` at line 12.

- If it returns `False` then $P_{\ell_P - 1} = \emptyset$ and it follows from Lemma 4.5 that $U_{\ell_P} \subseteq U_{\ell_P - 1}$. We get from the definition of $(U_k)_k$ that $U_k = U_{\ell_P - 1}$ for every $k \geq \ell_P$. Therefore, $U_{\ell_P - 1} = \bigcup_k U_k$. Moreover, $m_{init} \notin \uparrow B_{\ell_P - 1}$ because the condition of the **while** loop had to hold. It follows from Lemma 4.5 that $m_{init} \notin U_{\ell_P - 1}$. We derive from Lemma 3.2 that $m_{final} \notin Cov_{\mathcal{N}}$.
- If it returns `True` then $m_{init} \in \uparrow B_{\ell_B - 1}$ and it follows from Lemma 4.5 that $m_{init} \in U_{\ell_B - 1}$. We derive from Lemma 3.2 that $m_{final} \in Cov_{\mathcal{N}}$.

□

C Proofs for Section 6

Lemma 6.1. *We have* $\text{prop}_t(Q) \subseteq P \backslash Z$ *for every set* $Q \subseteq P \backslash Z$.

Proof. We can assume without loss of generality that $\bigwedge_{p \in P \backslash Q} F(p,t) = 0$ since otherwise the set $\text{prop}_t(Q)$ is empty. For the same reason, we can assume that there exists $q \in P$ such that $F(t,q) > 0$. Let us prove that such a place q cannot be in Z. We introduce the markings m_t and m'_t defined by $m_t(p) = F(p,t)$ and $m'_t(p) = F(t,p)$ for every $p \in P$. Those markings are the minimal ones satisfying $m_t \xrightarrow{t} m'_t$. Observe that for every $p \in Z$, we have $p \in P \backslash Q$ since $Q \cap Z$ is empty. It follows that $F(p,t) = 0$ for every $p \in Z$. Hence m_t is in the inductive invariant I_Z. Since $m_t \xrightarrow{t} m'_t$, we deduce that $m'_t \in I_Z$. As $F(t,q) > 0$, we get $m'_t(q) > 0$. Hence $q \notin Z$. □

Lemma 6.2. *We have* $Z = P \backslash Q$.

Proof. Since $Q_0 \subseteq P \backslash Z$, Lemma 6.1 shows by induction that $Q_k \subseteq P \backslash Z$ for every k. It follows that $Q \subseteq P \backslash Z$. We derive $Z \subseteq P \backslash Q$. The converse inclusion is obtained by proving that the set $M = \{m \in \mathbb{N}^P \mid \bigwedge_{p \in P \backslash Q} m(p) = 0\}$ is an inductive invariant. First of all, since $Q_0 \subseteq Q$, we deduce that $m_{init} \in M$. Now let us consider $m \in M$ and a transition t such that $m \xrightarrow{t} m'$ for some marking m'. Observe that $m(p) \geq F(p,t)$ for every $p \in P$. In particular, for $p \in P \backslash Q$, the equality $m(p) = 0$ implies $F(p,t) = 0$. Assume by contradiction that $m' \notin M$. In that case, there exists $q \in P \backslash Q$ such that $m'(q) > 0$. Since $m'(q) = m(q) + F(t,q) + F(q,t)$ and $m(q) = 0 = F(q,t)$, we deduce that $F(t,q) > 0$. Thus $q \in \text{prop}_t(Q)$. By definition of Q, we get $q \in Q$ and we obtain a contradiction. Thus M is an inductive invariant. By maximality of Z, we get the inclusion $P \backslash Q \subseteq Z$. Thus $Z = P \backslash Q$. □